

## CLAIMS

We claim:

1. A method for storing data, comprising:
  - 5 distributing a first plurality of groups of logical addresses among one or more storage devices;
  - receiving a second plurality of data-sets containing the data to be stored;
  - 10 assigning each data-set among the plurality of data-sets a number chosen from a first plurality of different numbers;
  - partitioning each data-set into multiple partitions, so that each partition among the multiple partitions receives a sequential partition number;
  - 15 assigning each partition within each data-set to be stored at a specific group of logical addresses in accordance with the sequential partition number of the partition and the number assigned to the data-set; and
  - storing each partition at the assigned specific group of logical addresses.
- 20 2. The method according to claim 1, wherein the multiple partitions comprise equal size partitions.
3. The method according to claim 1, wherein the data-sets comprise data from at least one of a file, file meta-data, a storage object, a data packet, a video tape, 25 a music track, an image, a database record, contents of a logical unit, and an email.
4. The method according to claim 1, wherein the first plurality of groups comprises  $s$  groups each having a different integral group number between 1 and  $s$ , wherein 30 the number comprises an integer  $r$  randomly chosen from and including integers between 0 and  $s-1$ , wherein the sequential partition number comprises a positive integer

p, and wherein the group number of the assigned specific group is  $(r+p)\bmodulo(s)$  if  $(r+p)\bmodulo(s)\neq 0$ , and s if  $(r+p)\bmodulo(s)=0$ .

5. The method according to claim 1, wherein the one or  
more storage devices are operative in at least one of a  
storage area network, a network attached storage system,  
and an object storage architecture.

6. The method according to claim 1, wherein the number  
is chosen by a randomizing function.

10 7. The method according to claim 1, wherein the number  
is chosen by a consistent hashing function.

8. A method for data distribution, comprising:  
receiving at least part of a data-set containing  
data;

15 delineating the data into multiple partitions;  
distributing logical addresses among an initial set  
of storage devices so as to provide a balanced access to  
the devices;

20 transferring the partitions to the storage devices  
in accordance with the logical addresses;

adding an additional storage device to the initial  
set, thus forming an extended set of the storage devices  
comprising the initial set and the additional storage  
device; and

25 redistributing the logical addresses among the  
storage devices in the extended set so as to cause a  
portion of the logical addresses and the partitions  
stored thereat to be transferred from the storage devices  
in the initial set to the additional storage device,  
30 while maintaining the balanced access and without  
requiring a substantial transfer of the logical addresses  
among the storage devices in the initial set.

9. The method according to claim 8, wherein the data-

set comprises data from at least one of a file, file meta-data, a storage object, a data packet, a video tape, a music track, an image, a database record, contents of a logical unit, and an email.

5 10. The method according to claim 8, wherein the initial set of storage devices and the additional storage device are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

10 11. The method according to claim 8, wherein distributing the logical addresses comprises:

generating a first plurality of sets of logical addresses,

and wherein delineating the data comprises:

15 assigning the at least part of the data-set a number chosen from a first plurality of different numbers; and

assigning each partition among the multiple partitions a sequential partition number,

and wherein transferring the partitions comprises:

20 storing each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number.

12. A method for data distribution, comprising:

receiving at least part of a data-set containing

25 data;

delineating the data into multiple partitions;

distributing logical addresses among an initial set of storage devices so as to provide a balanced access to the devices;

30 transferring the partitions to the storage devices in accordance with the logical addresses;

removing a surplus storage device from the initial set, thus forming a depleted set of the storage devices

comprising the initial set less the surplus storage device; and

5 redistributing the logical addresses among the storage devices in the depleted set so as to cause the logical addresses of the surplus device and the partitions stored thereat to be transferred to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the depleted set.

10 13. The method according to claim 12, wherein the data-set comprises data from at least one of a file, file meta-data, a storage object, a data packet, a video tape, a music track, an image, a database record, contents of a logical unit, and an email.

15 14. The method according to claim 12, wherein the initial set of storage devices is operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

15. The method according to claim 12, wherein  
20 distributing the logical addresses comprises:

generating a first plurality of sets of logical addresses,

and wherein delineating the data comprises:

25 assigning the at least part of the data-set a number chosen from a first plurality of different numbers; and

assigning each partition among the multiple partitions a sequential partition number,

and wherein transferring the partitions comprises:

30 storing each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number.

16. A data storage system, comprising:

one or more mass-storage devices, coupled to store

partitions of data at respective first ranges of logical addresses (LAs);

a plurality of interim devices, configured to operate independently of one another, each interim device being assigned a respective second range of the LAs and coupled to receive the partitions of data from and provide the partitions of data to the one or more mass-storage devices having LAs within the respective second range; and

one or more interfaces, which are adapted to receive input/output (IO) requests from host processors, to identify specified partitions of data in response to the IO requests, to convert the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data, and to direct all the converted-IO-requests to the interim device to which the specified LAs are assigned.

17. The storage system according to claim 16, wherein at least one of the mass-storage devices has a slow access time, and wherein at least one of the interim devices has a fast access time.

18. The storage system according to claim 16, wherein the one or more mass-storage devices are coupled to provide a balanced access to the first ranges of LAs.

25 19. The storage system according to claim 16, wherein the storage system is operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

20. A data storage system, comprising:

30 one or more storage devices wherein are distributed a first plurality of groups of logical addresses; and a processing unit which is adapted to: receive a second plurality of data-sets containing

the data to be stored,

assign each data-set among the plurality of data-sets a number chosen from a first plurality of different numbers,

5 partition each data-set into multiple partitions, so that each partition among the multiple partitions receives a sequential partition number,

10 assign each partition within each data-set to be stored at a specific group of logical addresses in the one or more storage devices in accordance with the sequential partition number of the partition and the number assigned to the data-set, and

15 store each partition in the one or more storage devices at the assigned specific group of logical addresses.

21. The storage system according to claim 20, wherein the multiple partitions comprise equal size partitions.

22. The storage system according to claim 20, wherein the data-sets comprise data from at least one of a file, 20 file meta-data, a storage object, a data packet, a video tape, a music track, an image, a database record, contents of a logical unit, and an email.

23. The storage system according to claim 20, wherein the first plurality of groups comprises s groups each 25 having a different integral group number between 1 and s, wherein the number comprises an integer r randomly chosen from and including integers between 0 and s-1, wherein the sequential partition number comprises a positive integer p, and wherein the group number of the assigned 30 specific group is  $(r+p)\bmod(s)$  if  $(r+p)\bmod(s)\neq 0$ , and s if  $(r+p)\bmod(s)=0$ .

24. The storage system according to claim 20, wherein the one or more storage devices and the processing unit

are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

25. The storage system according to claim 20, wherein  
5 the number is chosen by a randomizing function.

26. The storage system according to claim 20, wherein  
the number is chosen by a consistent hashing function.

27. Data distribution apparatus, comprising:

10 an initial set of storage devices among which are  
distributed logical addresses so as to provide a balanced  
access to the devices;

15 an additional storage device to the initial set,  
thus forming an extended set of the storage devices  
comprising the initial set and the additional storage  
device; and

20 a processor which is adapted to receive at least  
part of a data-set containing data, to delineate the data  
into multiple partitions, to transfer the partitions to  
the initial set of storage devices in accordance with the  
logical addresses, to redistribute the logical addresses  
among the storage devices in the extended set so as to  
cause a portion of the logical addresses and the  
partitions stored thereat to be transferred from the  
storage devices in the initial set to the additional  
25 storage device, while maintaining the balanced access and  
without requiring a substantial transfer of the logical  
addresses among the storage devices in the initial set.

28. The apparatus according to claim 27, wherein the  
data-set comprises data from at least one of a file, file  
30 meta-data, a storage object, a data packet, a video tape,  
a music track, an image, a database record, contents of a  
logical unit, and an email.

29. The apparatus according to claim 27, wherein the

initial set of storage devices and the additional storage device are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

5 30. The apparatus according to claim 27, wherein the logical addresses comprise a plurality of sets of logical addresses, and wherein the processor is adapted to:

assign the at least part of the data-set a number chosen from a plurality of different numbers,

10 assign each partition among the multiple partitions a sequential partition number, and

store each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number.

15 31. Data distribution apparatus, comprising:

an initial set of storage devices among which are distributed logical addresses so as to provide a balanced access to the devices;

20 a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set; and

25 a processor which is adapted to receive at least part of a data-set containing data, to delineate the data into multiple partitions, to transfer the partitions to the initial set of storage devices in accordance with the logical addresses, to redistribute the logical addresses and the partitions stored thereat of the surplus storage device among the storage devices in the depleted set while maintaining the balanced access and without 30 requiring a substantial transfer of the logical addresses among the storage devices in the depleted set.

32. The apparatus according to claim 31, wherein the data-set comprises data from at least one of a file, file

meta-data, a storage object, a data packet, a video tape, a music track, an image, a database record, contents of a logical unit, and an email.

33. The apparatus according to claim 31, wherein the 5 initial set of storage devices is operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

34. The apparatus according to claim 31, wherein the logical addresses comprise a plurality of sets of logical 10 addresses, and wherein the processor is adapted to:

assign the at least part of the data-set a number chosen from a plurality of different numbers,

assign each partition among the multiple partitions a sequential partition number, and

15 store each partition at one of the sets of logical addresses in accordance with the sequential partition number of the partition and the number.

35. A method for storing data, comprising:

coupling one or more mass-storage devices to store 20 partitions of data at respective first ranges of logical addresses (LAs);

configuring a plurality of interim devices to operate independently of one another;

25 assigning each interim device a respective second range of the LAs;

coupling each interim device to receive the partitions of data from and provide the partitions of data to the one or more mass-storage devices having LAs within the respective second range;

30 receiving input/output (IO) requests from host processors;

identifying specified partitions of data in response to the IO requests;

converting the IO requests to converted-IO-requests directed to specified LAs in response to the specified partitions of data; and

5 directing all the converted-IO-requests to the interim device to which the specified LAs are assigned.

36. The method according to claim 35, wherein at least one of the mass-storage devices has a slow access time, and wherein at least one of the interim devices has a fast access time.

10 37. The method according to claim 35, wherein the one or more mass-storage devices are coupled to provide a balanced access to the first ranges of LAs.

15 38. The method according to claim 35, wherein the one or more storage devices and the plurality of interim devices are operative in at least one of a storage area network, a network attached storage system, and an object storage architecture.

39. A method for data distribution, comprising:

20 receiving at least part of a data-set containing data;

delineating the data into multiple equal size partitions;

25 transferring the partitions to an initial set of storage devices so as to provide a balanced access to the devices;

adding an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device; and

30 redistributing the partitions among the storage devices in the extended set so as to cause a portion of the partitions to be transferred from the storage devices in the initial set to the additional storage device,

while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the initial set.

40. A method for data distribution, comprising:

5 receiving at least part of a data-set containing data;

delineating the data into multiple equal size partitions;

10 transferring the partitions to an initial set of storage devices so as to provide a balanced access to the devices;

15 removing a surplus storage device from the initial set, thus forming a depleted set of the storage devices comprising the initial set less the surplus storage device; and

20 redistributing the partitions stored in the surplus device to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the depleted set.

41. Data distribution apparatus, comprising:

an initial set of storage devices;

25 an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device; and

30 a processor which is adapted to receive at least part of a data-set containing data, to delineate the data into multiple equal size partitions, to transfer the partitions to the initial set of storage devices so as to provide a balanced access to the initial set of storage devices, to redistribute the partitions among the storage devices in the extended set so as to cause a portion the partitions stored in the initial set to be transferred to

the additional storage device, while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the initial set.

5 42. Data distribution apparatus, comprising:  
an initial set of storage devices;  
a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set; and  
10 a processor which is adapted to receive at least part of a data-set containing data, to delineate the data into multiple equal size partitions, to transfer the partitions to the initial set of storage devices so as to provide a balanced access to the initial set of storage  
15 devices, to redistribute the partitions of the surplus storage device among the storage devices in the depleted set while maintaining the balanced access and without requiring a substantial transfer of the partitions among the storage devices in the depleted set.